



US009208079B2

(12) **United States Patent**
Oh et al.

(10) **Patent No.:** **US 9,208,079 B2**
(45) **Date of Patent:** **Dec. 8, 2015**

(54) **SOLID STATE MEMORY (SSM), COMPUTER SYSTEM INCLUDING AN SSM, AND METHOD OF OPERATING AN SSM**

(71) Applicant: **SAMSUNG ELECTRONICS CO., LTD.**, Suwon-Si, Gyeonggi-Do (KR)

(72) Inventors: **Moon-Wook Oh**, Seoul (KR); **Do-Geun Kim**, Seoul (KR); **Chan-Ik Park**, Seoul (KR)

(73) Assignee: **Samsung Electronics Co., Ltd.**, Suwon-si, Gyeonggi-do (KR)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **14/717,476**

(22) Filed: **May 20, 2015**

(65) **Prior Publication Data**

US 2015/0261667 A1 Sep. 17, 2015

Related U.S. Application Data

(60) Continuation of application No. 14/445,126, filed on Jul. 29, 2014, and a continuation of application No. 14/146,870, filed on Jan. 3, 2014, and a division of application No. 14/101,469, filed on Dec. 10, 2013, which is a division of application No. 13/647,630, filed on Oct. 9, 2012, now Pat. No. 8,626,996, application No. 14/717,476, which is a division of application No. 13/027,299, filed on Feb. 15, 2011, now abandoned, and a division of application No. 12/015,548, filed on Jan. 17, 2008, now abandoned.

(30) **Foreign Application Priority Data**

Aug. 14, 2007 (KR) 10-2007-0081832

(51) **Int. Cl.**
G06F 12/02 (2006.01)
G11C 11/56 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 12/0246** (2013.01); **G11C 11/56** (2013.01); **G06F 2212/7202** (2013.01)

(58) **Field of Classification Search**
CPC G06F 12/0246; G11C 11/56
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,671,388 A	9/1997	Hasbun
5,737,742 A	4/1998	Achiwa et al.
5,802,554 A	9/1998	Caceres et al.
5,966,326 A	10/1999	Park et al.

(Continued)

FOREIGN PATENT DOCUMENTS

JP	06028246 A	2/1994
JP	2003216460	7/1998

(Continued)

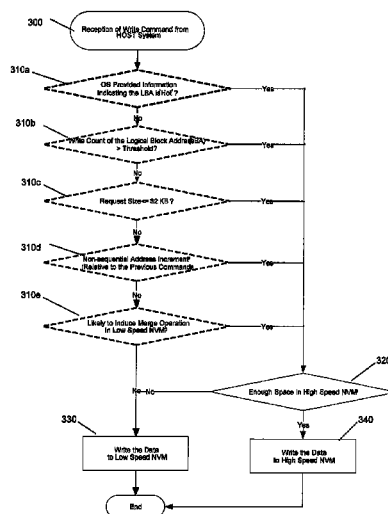
Primary Examiner — Hiep Nguyen

(74) *Attorney, Agent, or Firm* — Volentine & Whitt, PLLC

(57) **ABSTRACT**

In one aspect, data is stored in a solid state memory which includes first and second memory layers. A first assessment is executed to determine whether received data is hot data or cold data. Received data which is assessed as hot data during the first assessment is stored in the first memory layer, and received data which is first assessed as cold data during the first assessment is stored in the second memory layer. Further, a second assessment is executed to determine whether the data stored in the first memory layer is hot data or cold data. Data which is then assessed as cold data during the second assessment is migrated from the first memory layer to the second memory layer.

30 Claims, 8 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

6,330,572 B1 12/2001 Sitka
 6,701,421 B1 3/2004 Elnozahy et al.
 6,785,767 B2 8/2004 Coulson
 7,457,897 B1 11/2008 Lee et al.
 7,975,105 B1 * 7/2011 Sun et al. 711/115
 2002/0051394 A1 5/2002 Tobita et al.
 2003/0140207 A1 7/2003 Nagase et al.
 2005/0081009 A1 4/2005 Williams et al.
 2005/0251617 A1 11/2005 Sinclair et al.
 2005/0268062 A1 12/2005 Nagase et al.
 2006/0152981 A1 7/2006 Ryu
 2006/0248259 A1 11/2006 Ryu et al.
 2007/0011421 A1 1/2007 Keller et al.
 2007/0061502 A1 * 3/2007 Lasser et al. 711/103
 2007/0294474 A1 12/2007 Panabaker et al.
 2008/0016297 A1 1/2008 Bartley et al.
 2008/0104309 A1 * 5/2008 Cheon et al. 711/103
 2008/0126680 A1 * 5/2008 Lee et al. 711/103

2008/0209109 A1 8/2008 Lasser
 2008/0215800 A1 * 9/2008 Lee et al. 711/103
 2008/0215801 A1 9/2008 Tan et al.
 2008/0281939 A1 11/2008 Frazier et al.
 2009/0089486 A1 4/2009 Cheung et al.
 2009/0132778 A1 * 5/2009 Danilak 711/167
 2009/0144501 A2 6/2009 Yim et al.

FOREIGN PATENT DOCUMENTS

JP 2002140229 A 5/2002
 JP 2003-022687 1/2003
 JP 10187537 7/2003
 JP 2006302255 A 11/2006
 KR 100205240 B1 4/1999
 KR 1020030011229 A 2/2003
 KR 20070008403 1/2007
 KR 20070008403 A 1/2007
 KR 10-0704037 B1 3/2007
 WO WO2007084079 A1 7/2007

* cited by examiner

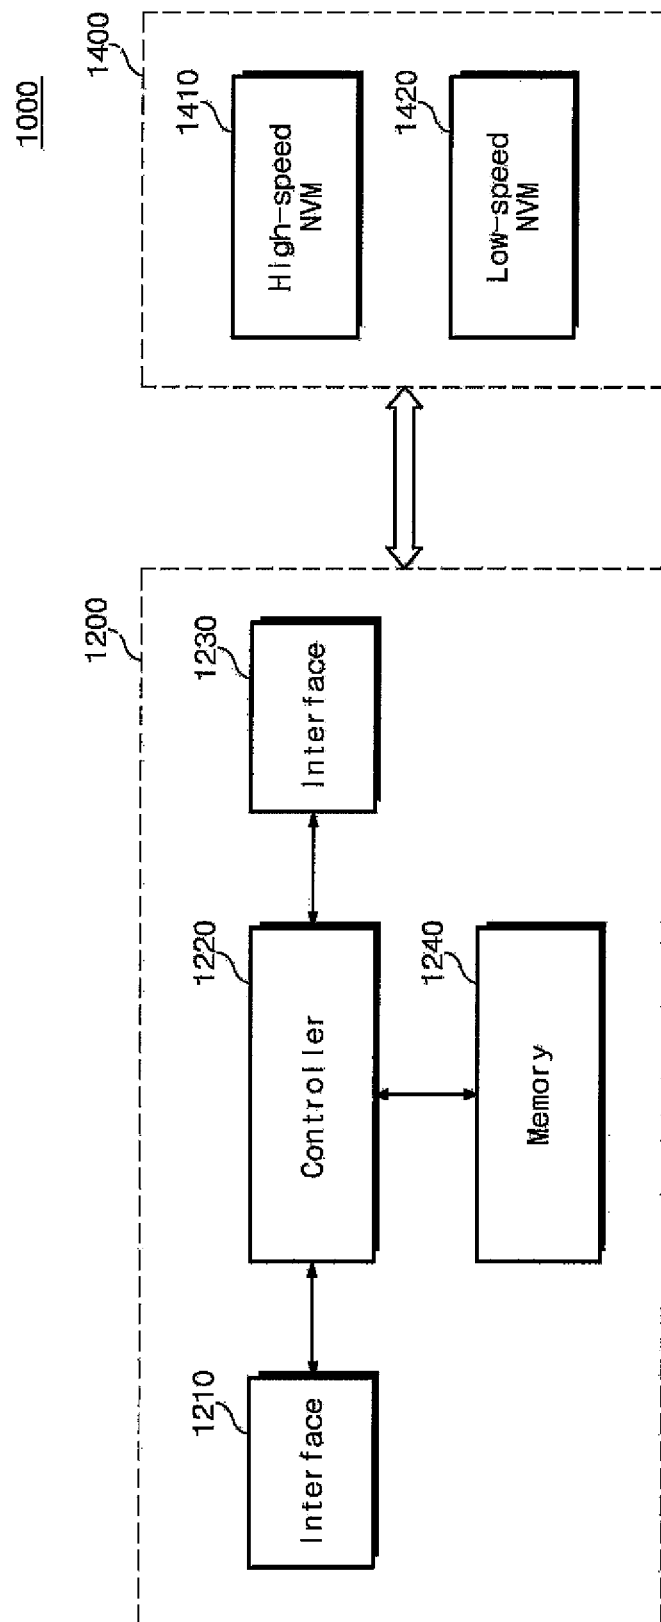


FIG. 1

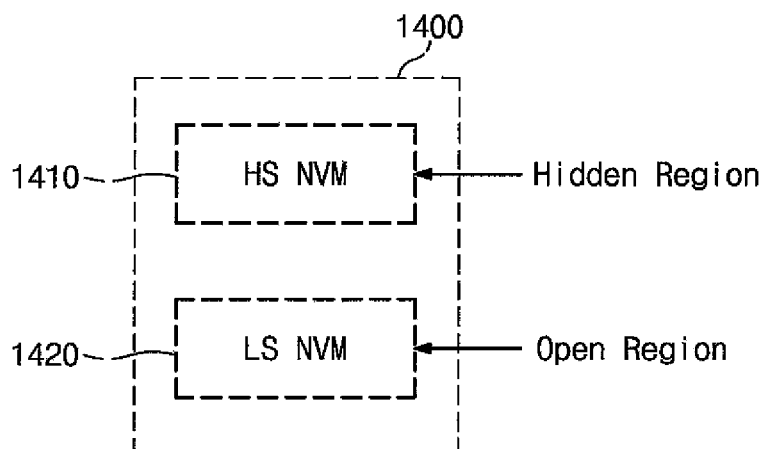


FIG. 2

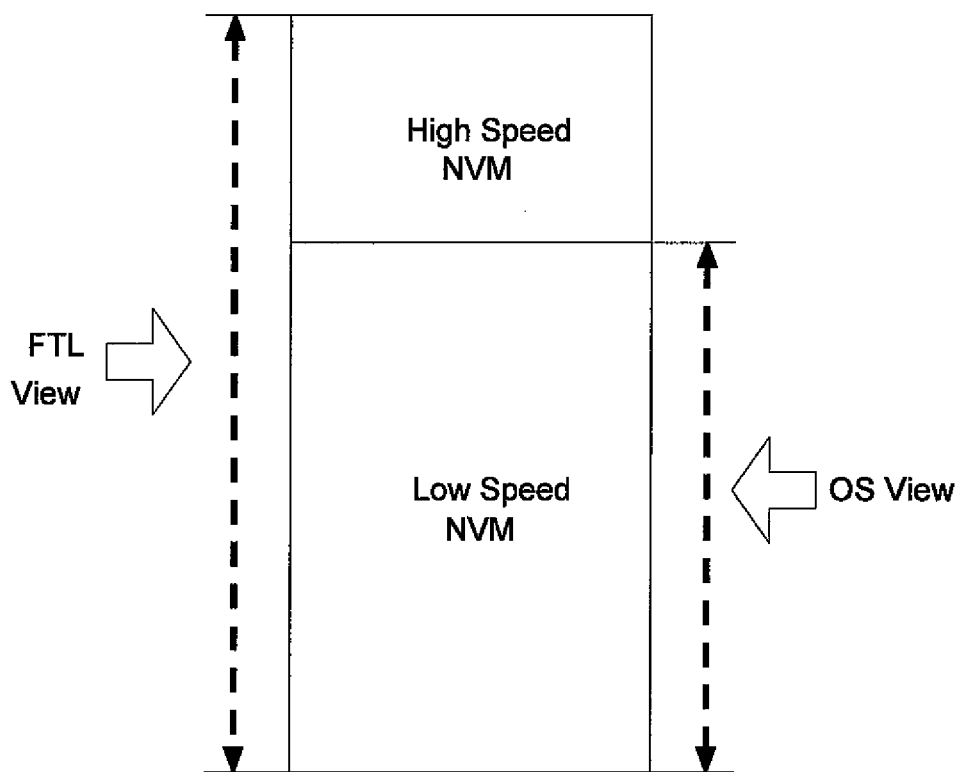


FIG. 3

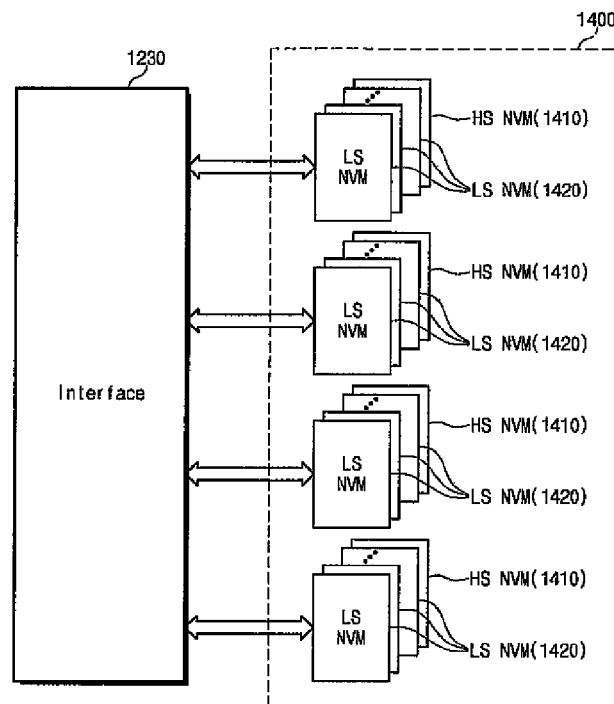


FIG. 4

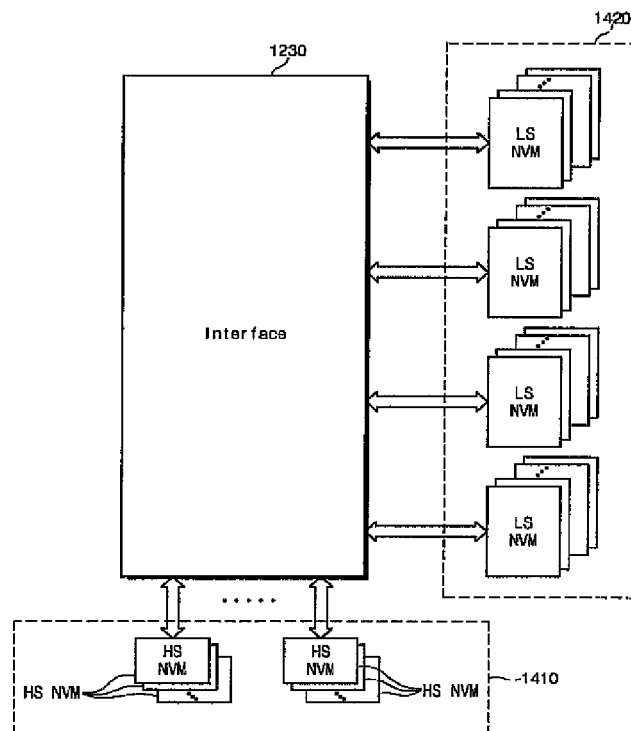
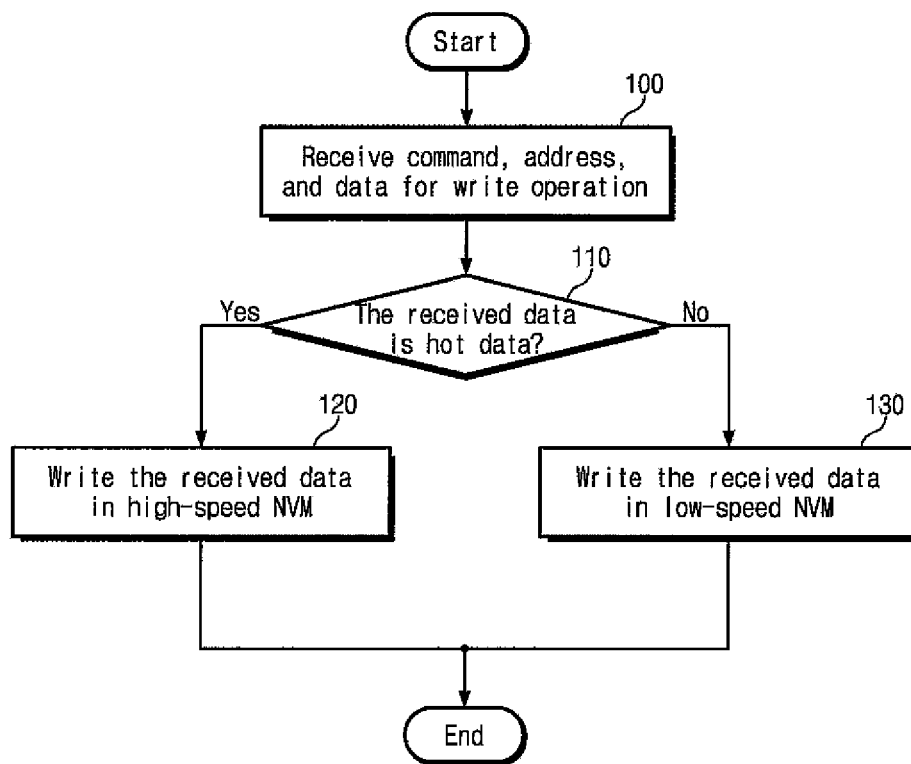


FIG. 5

**FIG. 6**

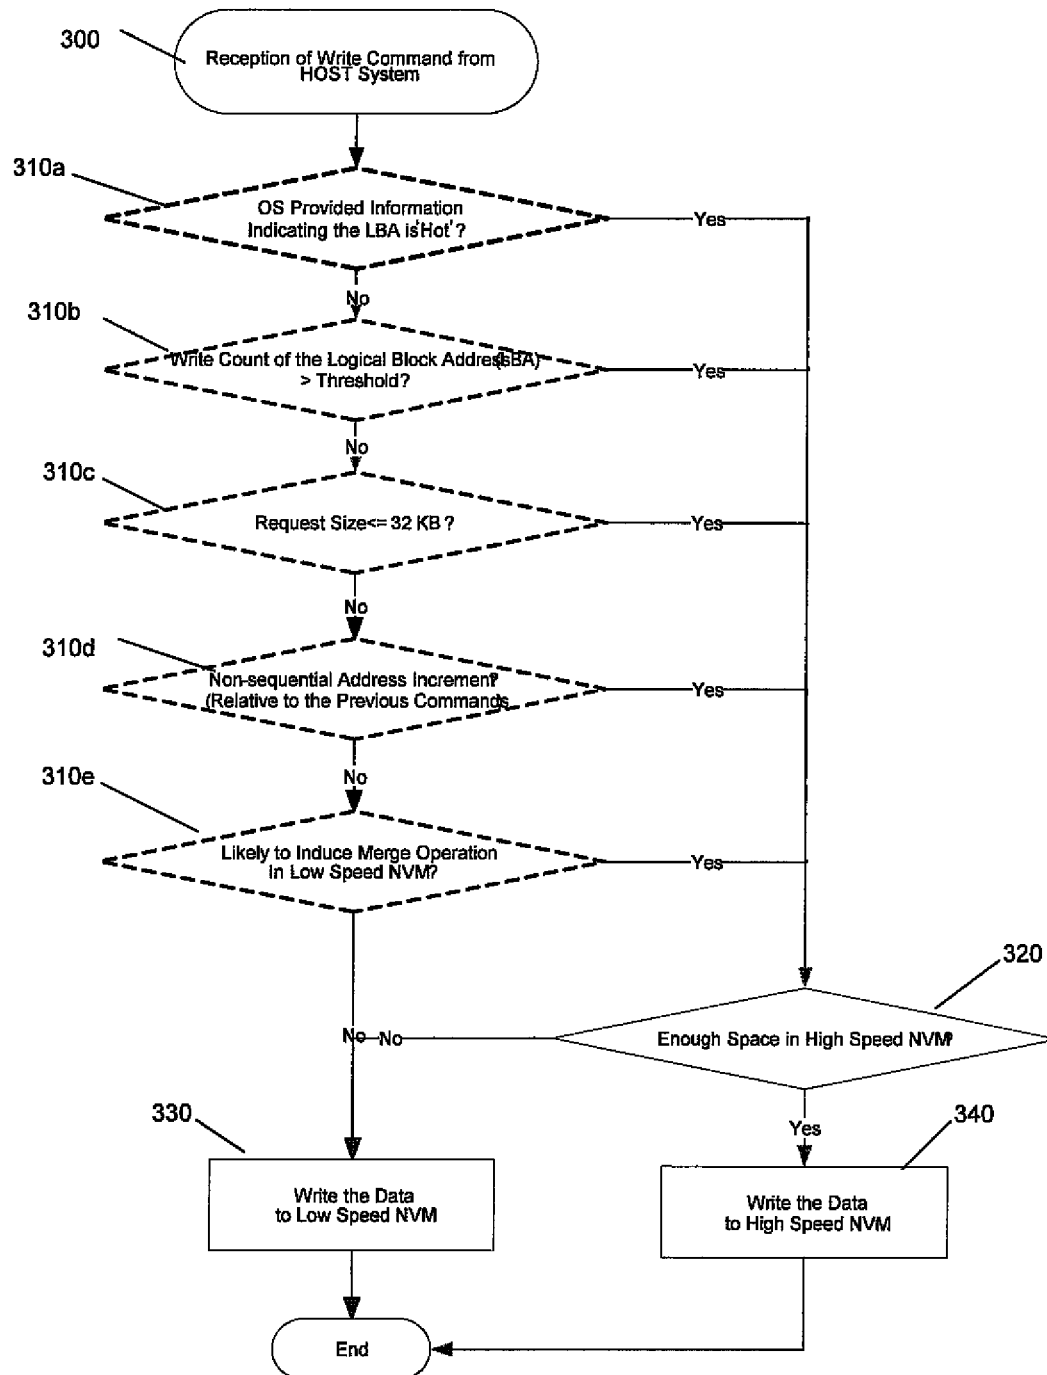


FIG. 7

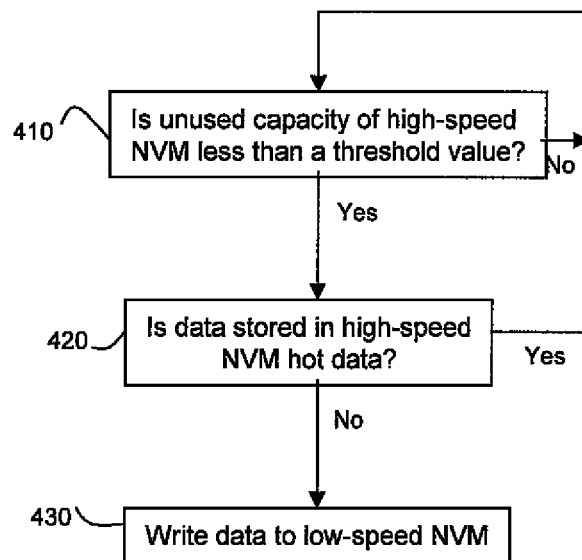
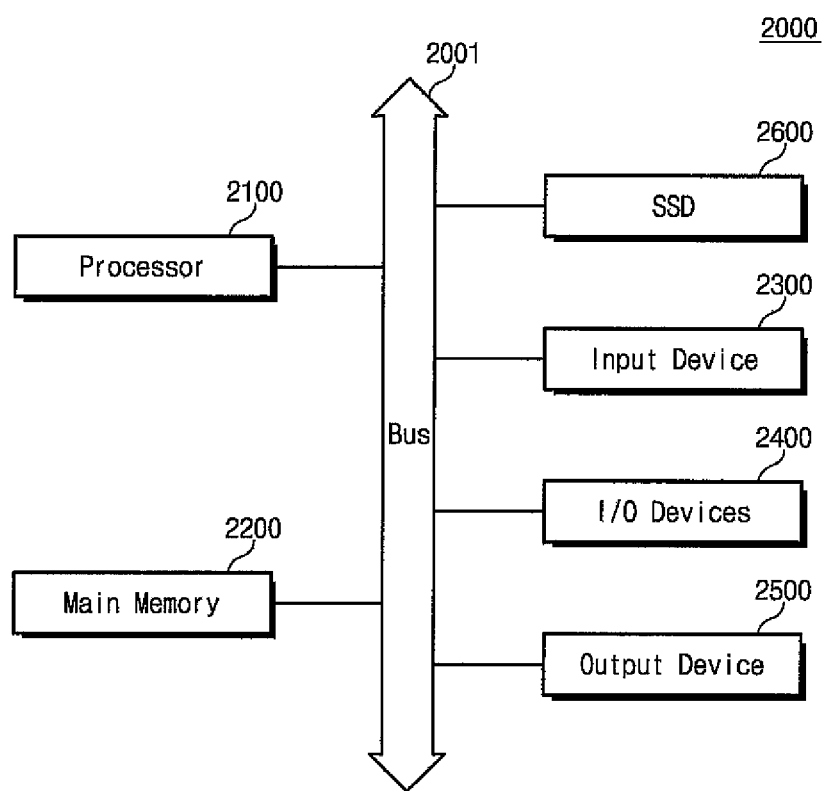
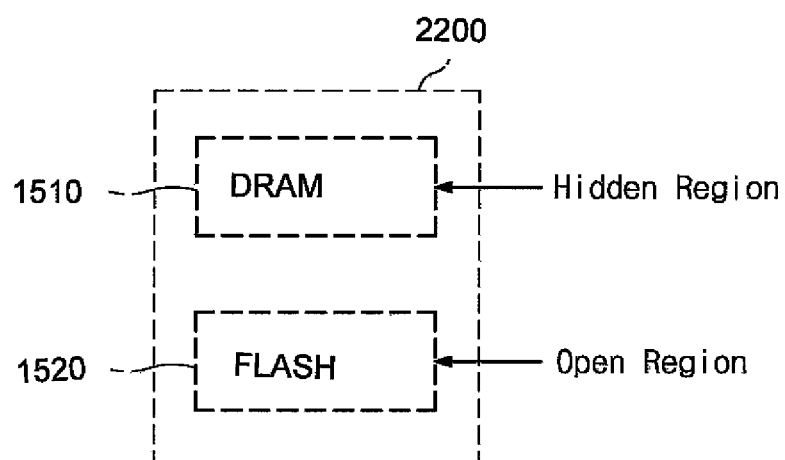
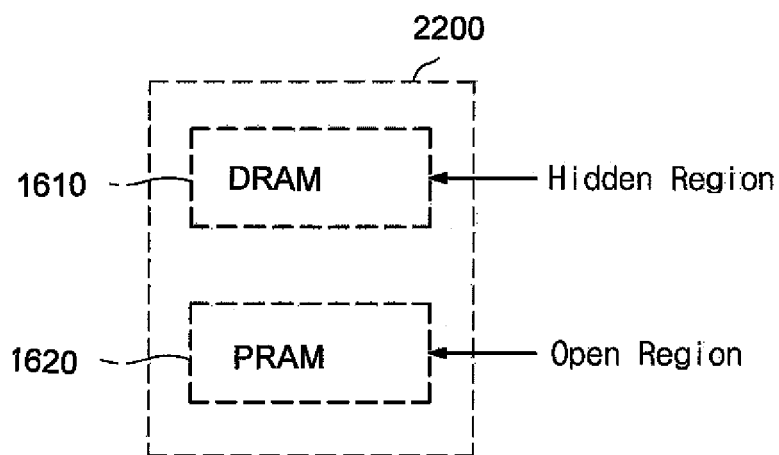


FIG. 8

**FIG. 9**

**FIG. 10****FIG. 11**

1

SOLID STATE MEMORY (SSM), COMPUTER SYSTEM INCLUDING AN SSM, AND METHOD OF OPERATING AN SSM

CROSS-REFERENCE TO RELATED APPLICATIONS

This is a continuation of co-pending application Ser. No. 14/445,126, filed Jul. 29, 2014, which is a continuation of co-pending application Ser. No. 14/146,870, filed Jan. 3, 2014, which is a continuation of co-pending application Ser. No. 14/101,469, filed Dec. 10, 2013, which is a divisional of application Ser. No. 13/647,630, filed Oct. 9, 2012, which issued as U.S. Pat. No. 8,626,996, on Jan. 7, 2014, and which is a divisional of abandoned application Ser. No. 13/027,299, filed Feb. 15, 2011, which is a divisional of abandoned application Ser. No. 12/015,548, filed Jan. 17, 2008, which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention generally relates to memory systems, and more particularly, the present invention relates to a solid state memory (SSM), a computer system which includes an SSM, and a method of operating an SSM. Examples of the SSM include the main memory of a computer system and the solid state drive (SSD) of a computer system.

A claim of priority is made to Korean Patent Application No. 2007-0081832, filed Aug. 14, 2007, the entirety of which is incorporated herein by reference.

2. Description of the Related Art

A solid state drive (SSD) is a data storage device that typically emulates a conventional hard disk drive (HDD), thus easily replacing the HDD in most applications. In contrast to the rotating disk medium of an HDD, an SSD utilizes solid state memory to store data. With no moving parts, an SSD largely eliminates seek time, latency and other electro-mechanical delays and failures associated with a conventional HDD.

An SSD is commonly composed of either NAND flash (non-volatile) or SDRAM (volatile).

SSDs based on volatile memory such as SDRAM are characterized by fast data access and are used primarily to accelerate applications that would otherwise be held back by the latency of disk drives. The volatile memory of the DRAM-based SSDs typically requires the inclusion of an internal battery and a backup disk system to ensure data persistence. If power is lost, the battery maintains power for sufficient duration of copy data from the SDRAM to the backup disk system. Upon restoration of power, data is copied back from the backup disk to SDRAM, at which time the SSD resumes normal operations.

However, most SSD manufacturers use non-volatile flash memory to create more rugged and compact alternatives to DRAM-based SSDs. These flash memory-based SSDs, also known as flash drives, do not require batteries, allowing makers to more easily replicate standard hard disk drives. In addition, non-volatile flash SSDs retain memory during power loss.

As is well known in the art, single-level cell (SLC) flash is capable of storing one bit per memory cell, while multi-level cell (MLC) flash is capable of storing two or more bits per memory cell. As such, in order to increase capacity, flash SSDs may utilize multi-level cell (MLC) memory banks. However, flash SSDs generally suffer from relatively slow random write speeds, and this operational drawback is further

2

exasperated with relatively slow speeds of MLC flash. As such, it has been suggested to equip SSDs with two types of flash storage media—lower capacity SLC memory banks and higher capacity MLC memory banks. With such a configuration, frequently used data (e.g., directory and/or log information) can be stored in the faster SLC banks, while less frequently used data (e.g., music files, images, etc.) can be stored in the slower MLC banks.

SUMMARY OF THE INVENTION

According to an aspect of the present invention, data is stored in a solid state memory which includes first and second memory layers. A first assessment is executed to determine whether received data is hot data or cold data. Received data which is assessed as hot data during the first assessment is stored in the first memory layer, and received data which is assessed as cold data during the first assessment is stored in the second memory layer. Further, a second assessment is executed to determine whether the data stored in the first memory layer is hot data or cold data. Data which is then assessed as cold data during the second assessment is migrated from the first memory layer to the second memory layer.

According to another aspect of the present invention, a method of storing received data in a solid state memory includes initially storing hot data in a high-speed memory layer and, and then migrating a portion of the data stored in the high-speed memory layer to a low-speed memory layer for storing cold data.

According to yet another aspect of the present invention, a solid state memory system includes a first memory layer, a second memory layer, and a memory controller. The memory controller is configured to execute a first assessment of whether received data is hot data or cold data, to store received data which is assessed as hot data during the first assessment in the first memory layer, and to store received data which is assessed as cold data during the first assessment in the second memory layer. The memory controller is further configured to execute a second assessment of whether the data stored in the first memory layer is hot data or cold data, and to migrate data which is assessed as cold data during the second assessment from the first memory layer to the second memory layer.

According to still another aspect of the present invention, a solid state memory system is configured to operatively connect to a computer operating system and comprises first and second memory layers. An operational speed of the first memory layer is greater than an operational speed of the second memory layer, and the first memory area is operationally hidden from the computer operating system when the solid state memory is operatively connected to the computer operating system.

According to another aspect of the present invention, a computer system includes a processor and a memory. The solid state memory includes a high-speed memory layer and a low-speed memory layer, and the high-speed memory area is operationally hidden from the processor.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other aspects and features of the present invention will become readily apparent from the detailed description that follows, with reference to the accompanying drawings, in which:

FIG. 1 is a block diagram of a solid state drive (SSD) according to an embodiment of the present invention;

FIGS. 2 and 3 are block diagrams for describing a non-volatile storage media in the SSD of FIG. 1 according to an embodiment of the present invention;

FIGS. 4 and 5 are block diagrams for describing alternative ways of coupling a non-volatile storage media to an interface in the SSD of FIG. 1 according to embodiments of the present invention;

FIGS. 6 through 8 are flow charts for use in describing methods of allocating data to regions of a non-volatile storage media in the SSD of FIG. 1 according to embodiments of the present invention;

FIG. 9 is a block diagram of the computer system including an SSD according to an embodiment of the present invention; and

FIGS. 10 and 11 are block diagrams of a main memory according to embodiments of the present invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The present invention will now be described by way of preferred, but non-limiting, embodiments of the invention. It is emphasized here that the invention is not limited by the exemplary embodiments described below, and that instead the scope of the invention is delimited by the appended claims.

FIG. 1 illustrates a block diagram of a solid state drive (SSD) 1000 according to an embodiment of the present invention. As shown, the SSD 1000 of this example includes an SSD controller 1200 and non-volatile storage media 1400.

The SSD controller 1200 includes first and second interfaces 1210 and 1230, a controller 1220, and a memory 1240.

The first interface 1210 functions as a data I/O interface with a host device, such as a host central processing unit (CPU) (not shown). Non-limiting examples of the first interface 1210 include Universal Serial Bus (USB) interfaces, Advanced Technology Attachment (ATA) interfaces, Serial ATA (SATA) interfaces, Small Computer System Interface (SCSI) interfaces.

The second interface 1230 functions as a data I/O interface with the non-volatile storage media 1400. In particular, the second interface 1230 is utilized to transmit/receive various commands, addresses and data to/from the non-volatile storage media 1400. As will be apparent to those skilled in the art, a variety of different structures and configurations of the second interface 1230 are possible, and thus a detailed description thereof is omitted here for brevity.

The controller 1220 and memory 1240 are operatively connected between the first and second interfaces 1210 and 1230, and together function to control/manage the flow of data between the host device (not shown) and the non-volatile storage media 1400. The memory 1240 may, for example, be an DRAM type of memory device, and the controller 1220 may, for example, include a central processing unit (CPU), a direct memory access (DMA) controller, and an error correction control (ECC) engine. Examples of controller functionality may be found in commonly assigned U.S. Patent Publication 2006-0152981, which is incorporated herein by reference. The operations generally executed by controller 1220 (and memory 1240) to transfer data between the host device (not shown) and SSD memory banks are understood by those skilled in the art, and thus a detailed description thereof is omitted here for brevity. Rather, the operational description presented later herein is primarily focused on inventive aspects relating to various embodiments of the invention.

Still referring to FIG. 1, the non-volatile storage media 1400 of this example includes a high-speed non-volatile memory (NVM) 1410 and a low-speed non-volatile memory (NVM) 1420. As the names suggest, the high-speed NVM 1410 is capable of operating at a relatively higher speed (e.g., random write speed) when compared to the low-speed NVM 1420.

In an exemplary embodiment, the high-speed NVM 1410 is single-level cell (SLC) flash memory, and the low-speed NVM 1420 is multi-level cell (MLC) flash memory. However, the invention is not limited in this respect. For example, the high-speed NVM 1410 may instead be comprised of phase-change random access memory (PRAM), or MLC flash memory in which one bit per cell is utilized. Also, the high-speed NVM 1410 and the low-speed NVM 1420 may be comprised of the same type of memory (e.g., SLC or MLC or PRAM), where the operational speed is differentiated by fine-grain mapping in the high-speed NVM 1410 and coarse-grain mapping in the low-speed NVM 1420.

Generally, the high-speed NVM 1410 is utilized to store frequently accessed (written) data such as meta data, and the low-speed NVM 1420 is utilized to store less frequently accessed (written) data such as media data. In other words, as will discussed later herein, a write frequency of data in the high-speed NVM 1410 is statistically higher than a write frequency of data in the low-speed NVM 1420. Also, due to the nature of the respective data being stored, the storage capacity of the low-speed NVM 1420 will typically be much higher than that of the high-speed NVM 1410.

In an exemplary embodiment, the high-speed NVM 1410 is hidden from an external operating system connected to the SSD. This aspect of the embodiment is illustrated in FIGS. 2 and 3.

Referring collectively to FIGS. 2 and 3, the high-speed NVM 1410 of the example of this embodiment is a hidden region—that is, the high-speed NVM 1410 is cannot be seen (directly addressed) by the external operating system (OS). Rather, the address space shown relative to the OS view is only the low-speed NVM 1420. On the other hand, the address space shown relative to the Flash Translation Layer (FTL) is both the high-speed NVM 1410 and the low-speed NVM 1420. The FTL translates an address provided by the OS into a physical address of the non-volatile storage media 1400 (i.e., a physical address within the high-speed NVM 1410 or the low-speed NVM 1420).

Turning to the block diagrams of FIGS. 4 and 5, there are a number of different ways in which the high-speed NVM 1410 and the low-speed NVM 1420 can be operative connected to the controller 1220 (FIG. 1) via the interface 1230. In the example of FIG. 4, the high-speed NVM 1410 and the low-speed NVM 1420 communicate via the interface 1230 using common interface channels. In the example of FIG. 5, the high-speed NVM 1410 and the low-speed NVM 1420 communicate via the interface 1230 using separate interface channels.

It is again noted, however, that the high-speed NVM 1410 and the low-speed NVM 1420 need not be composed of different types of memory. That is, a single type of memory may be operationally segregated into a high-speed layer and a low-speed layer. For example, the grain mapping in the two layers may differ, or the number of bits utilized per cell in the two layers may differ. Further, the high-speed memory layer and the low-speed memory layer may be segregated at the chip level (e.g., contained in different memory chips), or within the same memory chip (e.g., contained in different memory blocks or groups of memory cells of the same memory chip).

5

An operational description of the SSD according to embodiments of the present invention is presented next.

According to an embodiment of the present invention in which data is stored in the SSD, a first assessment is executed to determine whether received data is hot data or cold data. As will be understood by those skilled in the art, "hot" data is a term of art that refers to data which is frequently written or updated (requiring write access), such as directory information and/or logging information. "Cold" data is all other data, i.e., data which is not frequently written or updated, such as image files, sound files, program code and so on. Cold data may be written once or infrequently, but read frequently. Thus, it is the frequency of write access that separates hot data from cold.

Received data which is assessed as hot data during the first assessment is stored in the high-speed NVM **1410**, and received data which is first assessed as cold data during the first assessment is stored in the low-speed NVM **1420**.

Then, a second assessment is executed to determine whether the data stored in the high-speed NVM **1410** is hot data or cold data. In other words, the data stored in the high-speed NVM **1410** reassessed to determine with the data should be reclassified as cold data. Data which is then assessed as cold data during the second assessment is migrated from the high-speed NVM **1410** to the low-speed NVM **1420**.

By periodically migrating data which initially determined to be hot data from the high-speed NVM **1410** to the low-speed NVM **1420**, the size of the high-speed NVM **1410** can be reduced. This can potentially result in cost savings, and increase the overall storage capacity of the SSD (e.g., by allowing for more space for the high-capacity MLC layer).

The second assessment and migration of data to the low-speed NVM **1420** can be programmed to occur, for example, when the unused capacity of the high-speed NVM **1410** is less than a preset value. Alternately, for example, the second assessment and migration of data to the low-speed NVM **1420** can be programmed to occur at given periodic intervals, or when the SSD is idle. Examples of an idle state may include periods in which no read/write request is received from the host, or when the activation ratio or intensity of read/write requests is less than a threshold.

FIG. 6 is a flow chart for use in describing the first assessment and storage (write) of data in the SSD according to an embodiment of the present invention.

Initially, at step **100**, a write command, an address and data are received. Then, at step **110**, a determination is made as to whether the received data is classified as hot data. If the received data is classified as hot data, the received data is stored in the high-speed NVM **1410** at step **120**. On the other hand, if the received data is not classified as hot data, the received data is stored in the low-speed NVM **1420** at step **130**.

It should be noted that data stored in the low-speed NVM **1420** at step **130** may first be "passed through" the high-speed NVM **1410**. In other words, the data may first be briefly (temporarily) stored in the high-speed NVM **1410**, and then stored in the low-speed NVM. In this case, the high-speed NVM **1410** essentially acts as a memory buffer for the low-speed NVM **1420**.

FIG. 7 is a flow chart for use in describing the first assessment and storage (write) of data in the SSD according to another embodiment of the present invention.

Initially, at step **300**, a write command, an address and data are received. Then, at steps **310a** through **310e**, a determination is made as to whether the received data is to be classified as hot data. If the received data is classified as hot data, and if

6

there is sufficient available space in the high-speed NVM **1410** (step **320**), the received data is stored in the high-speed NVM **1410** at step **340**. On the other hand, if the received data is not to be classified as hot data, or if there is insufficient available space in the high-speed NVM **1410**, the received data is stored in the low-speed NVM **1420** at step **330**.

There are a number of different ways in which the received data might be classified as hot data, and steps **310a** through **310e** of FIG. 7 represent a non-exhaustive list of decision processes which can be used in the classification. These steps can be used in combinations of two or more, or individually, depending on the desired level of accuracy in the first assessment of the received data.

At step **310a**, a determination is made as to whether the operating system (OS) has provided information that the data is hot data. If so, the data is classified as hot data, and the process proceeds to step **320**.

At step **310b**, a determination is made as to whether the write count of the logical block address has exceeded a predetermined threshold. If so, the data is classified as hot data, and the process proceeds to step **320**.

At step **310c**, a determination is made as to whether the request size of the data is less than predetermined threshold (e.g., less than 32 KB). If so, the data is classified as hot data, and the process proceeds to step **320**.

At step **310d**, a determination is made as to whether there is a non-sequential address increment relative to the previously received command. If so, the data is classified as hot data, and the process proceeds to step **320**.

At step **310e**, a determination is made as to whether a merge operation is likely to be induced in the low-speed NVM. If so, the data is classified as hot data, and the process proceeds to step **320**.

Although not shown in FIG. 7, in the case where insufficient space exists in the high-speed NVM (step **320**), and alternative would be to create available space by migrating already stored cold data of the high-speed NVM to the low-speed NVM, and then storing the new hot data in the high-speed NVM.

Also, with reference to above-described processes of FIGS. 6 and 7, it is noted that the embodiments thereof are not limited to storing all of the hot and cold data in the high-speed and low-speed NVMs, respectively. For example, some of the data initially assessed as cold data may be stored in the high-speed NVM. Also, though less preferable, some of the data initially assessed as hot data may be stored in the low-speed NVM.

FIG. 8 is a flow chart for use in describing an example of the second assessment and migration of data to the low-speed NVM in the SSD according to an embodiment of the present invention.

Initially, at step **410**, a determination is made as to whether an unused memory capacity of the high-speed NVM is less than a predetermined threshold value. As suggested previously, this step can be supplemented with (or replaced with) a periodic execution step in which step **410** (or step **420** below) is executed at periodic intervals, and/or with a SSD idle determination step in which step **410** (or step **420**) is executed at periodic intervals.

Next, at step **420**, a determination is made as to whether data stored in the high-speed NVM is hot data, i.e., whether the data may be reclassified as cold data.

Then, at step **430**, reclassified cold data which stored in the high-speed NVM is migrated to the low-speed NVM.

There are a number of different ways in which the determination of step **420** may be executed. For example, it is possible to examine the write count value of each valid data in

7

the high-speed NVM, and to then reclassify data having low write counts as cold data. Alternately, it is possible to carry out a FIFO-type assessment in which old (first come) valid data is reclassified as cold data.

With reference to above-described process of FIG. 8, it is noted that the embodiment thereof is not limited to migrating all of the cold data to the low-speed NVM. For example, some of the data assessed as cold data may be retained in the high-speed NVM.

FIG. 9 is a block diagram of a computer system according to an embodiment of the present invention. As shown in the figure, a processor (host) **2100** and main memory **2200** communicate over a data bus **2001**. Also connected to the bus **2001** are an output device **2500** (e.g., display), an input device **2300** (e.g., keyboard), other I/O devices **2400**, and a solid state drive SSD. The solid state drive is configured according to one or more of the previously described embodiments of the invention.

Embodiments of the present invention have been described primarily in the context of solid state drives (SSDs). However, the invention is not limited to SSD applications. For example, FIG. 10 illustrates an embodiment where the high-speed memory layer and the low speed memory layer constitute the main memory **2200** of the computer system shown in FIG. 9. In FIG. 10, the high-speed memory layer **1510** includes DRAM cells and may be a hidden region relative to the processor **2100** of FIG. 9. The low-speed memory layer **1520** of FIG. 10 includes flash cells (either SLC or MLC) and may be open relative to the processor **2100** of FIG. 9. FIG. 11 illustrates another example of a main memory **2200**. As shown, the high-speed memory layer **1610** includes DRAM cells and may be a hidden region relative to the processor **2100** of FIG. 9, and the low-speed memory layer **1620** includes phase-change random access memory (PRAM) cells and may be open relative to the processor **2100** of FIG. 9.

The above-disclosed subject matter is to be considered illustrative, and not restrictive, and the appended claims are intended to cover all such modifications, enhancements, and other embodiments, which fall within the true spirit and scope of the present invention. Thus, to the maximum extent allowed by law, the scope of the present invention is to be determined by the broadest permissible interpretation of the following claims and their equivalents, and shall not be restricted or limited by the foregoing detailed description.

What is claimed is:

1. A memory system, comprising:

a first group of flash memory cells, each of the first group of flash memory cells configured to store a first number of bits per cell;

a second group of flash memory cells, each of the second group of flash memory cells configured to store a second number of bits per cell; and

a controller in communication with the first group of flash memory cells and the second group of flash memory cells, wherein the controller is configured to receive input data from a host and write the input data to one of the first group of flash memory cells and the second group of flash memory cells based on the size of the input data, and the controller is further configured to migrate first data from the first group of flash memory cells to the second group of flash memory cells if the number of available memory cells in the first group of flash memory cell is less than a first threshold.

2. The memory system of claim 1, wherein the first group of memory cells are single-level flash memory cells (SLC), and the second group of memory cells are multi-level flash memory cells (MLC).

8

3. The memory system of claim 2, wherein the first group of flash memory cells and the second group of flash memory cells are within a memory chip.

4. The memory system of claim 2, wherein the first data are cold data which are not frequently written or updated such as image files, sound files, or program codes.

5. The memory system of claim 2, wherein the controller is configured to write the input data to the second group of flash memory cells if the size of the input data is greater than a second threshold and to write the input data to first group of flash memory cells if the size of the input data is equal or less than the second threshold.

6. The memory system of claim 5, wherein the writing the input data includes determining whether the size of the input data is greater than the second threshold.

7. The memory system of claim 6, wherein the second threshold is 32 Kbyte.

8. The memory system of claim 2, wherein the controller is configured to determine whether the number of available memory cells in the first group of flash memory cell is less than the first threshold.

9. The memory system of claim 8, wherein the determining whether the number of available memory cells in the first group of flash memory cell is less than the first threshold occurs periodically with predetermined interval.

10. The memory system of claim 2, wherein the controller is further configured to select the first data among data stored in the first group of memory cells based on a write access frequency of the first data, and migrate the first data from the first group of memory cells to the second group of memory cells.

11. The memory system of claim 10, wherein the controller is further configured to manage a logical block address access for the selection of the first data.

12. The memory system of claim 2, wherein the first group of memory cells use a fine-grain mapping and the second group of memory cells use a coarse-grain mapping such that an operational speed of the first group of memory cells is faster than an operational speed of the second group of memory cells.

13. A method of storing data to a solid-state drive including a first group of flash memory cells, each of the first group of flash memory cells storing a first number of bits per cell, a second group of memory flash memory cells, each of the second group of flash memory cells storing a second number of bits per cell, a controller in communication with the first group of flash memory cells and the second group of flash memory cells, and a volatile memory device, the method comprising:

receiving input data from a host and storing the input data in the volatile memory device via the controller;

writing the input data to one of the first group of flash memory cells and the second group of flash memory cells via the controller based on the size of the input data; and

migrating first data from the first group of flash memory cells to the second group of flash memory cells if the number of available memory cells in the first group of flash memory cells is less than a first threshold.

14. The method of claim 13, wherein the writing the input data is performed to the second group of flash memory cells if the size of the input data is greater than a second threshold and to the first group of flash memory cells if the size of the input data is equal or less than the second threshold.

15. The method of claim 14, wherein the writing the input data further includes determining whether the size of the input data is greater than the second threshold.

16. The method of claim 15, wherein the second threshold is 32 Kbyte.

17. The method of claim 13, wherein the migrating occurs upon periodical determining whether the number of available memory cells of the first group of flash memory cells is less than the first threshold.

18. A solid state drive, comprising:

a first group of flash memory cells, each of the first group of flash memory cells configured to store 1-bit per cell;

a second group of flash memory cells, each of the second group of flash memory cells configured to store 2-bits per cell;

a controller configured to receive input data from a host via a first interface, and to write the input data, via a second interface, to one of the first group of flash memory cells and the second group of flash memory cells based on the size of the input data, and the controller is further configured to migrate first data from the first group of flash memory cells to the second group of flash memory cells if the number of available memory cells in the first group of flash memory cells is less than a first threshold; and a volatile memory device in communication with the controller via a third interface to store at least a portion of the input data temporarily before the input data is written to one of the first group of flash memory cells and the second group of flash memory cells.

19. The method of claim 18, wherein the volatile memory device is a DRAM.

20. The solid state drive of claim 18, wherein the first data are cold data which are not frequently written or updated such as image files, sound files, or program codes.

21. The solid state drive of claim 18, wherein the first interface is one of Universal Serial Bus (USB) interfaces, Advanced Technology Attachment (ATA) interfaces, Serial ATA (SATA) interfaces, and Small Computer System Interface (SCSI) interfaces.

22. The solid state drive of claim 18, wherein the second interface comprises a plurality of channels, each of the channels coupled to one of a plurality of nonvolatile storage media, one of the plurality of error checking and correcting ECC circuits and one of a plurality of buffer circuits.

23. The solid state drive of claim 18, wherein the controller further including a plurality ECC circuits configured to detect and correct errors on the write/read data transferred between the controller and the nonvolatile storage media via the second interface.

24. The solid state drive of claim 23, wherein the second interface comprises a plurality of channels, each of the channels coupled to one of a plurality of nonvolatile storage media, one of the plurality of error checking and correcting ECC circuits and one of a plurality of buffer circuits.

25. The solid state drive of claim 18, wherein the first group of flash memory cells and the second group of flash memory cells are within a memory chip.

26. The solid state drive of claim 18, wherein the controller is configured to write the input data to the second group of flash memory cells if the size of the input data is greater than a second threshold and to write the input data to the first group of flash memory cells if the size of the input data is equal or less than the second threshold.

27. The solid state drive of claim 26, wherein the writing the input data includes determining whether the size of the input data is greater than the second threshold.

28. The solid state drive of claim 27, wherein the second threshold is 32 Kbyte.

29. A solid state drive, comprising:

a nonvolatile storage media including a first group of flash memory cells, each of the first group of flash memory cells configured to store a first number of bits per cell, and a second group of flash memory cells, each of the second group of flash memory cells configured to store a second number of bits per cell;

a controller configured to:

receive input data from a host via a first interface; produce ECC data by using at least a portion of the input data;

write the input data and the ECC data to one of the first group of flash memory cells and the second group of flash memory cells via a second interface based on the size of the input data;

read output data from one of the first group of flash memory cells and the second group of flash memory cells via the second interface;

migrate first data from the first group of flash memory cells to the second group of flash memory cells if the number of available memory cells in the first group of flash memory cells is less than a first threshold; and

a DRAM in communication with the controller via a third interface to store at least a portion of the input data temporarily before the input data is written to the nonvolatile storage media,

wherein the controller includes a plurality of ECC circuits to detect and correct errors of the output data.

30. The solid state drive of claim 29, wherein the second interface comprises a plurality of channels, one of the plurality of ECC circuits dedicated to one of the plurality of channels, at least one of a plurality of nonvolatile storage media dedicated to the one of the plurality of channels, and at least one of a plurality of buffer circuits dedicated to the one of the plurality of channels.

* * * * *